arXiv:1611.03099v1 [cs.OH] 9 Nov 2016

A Brief Survey of Non-Residue Based Computational Error Correction

Sriseshan Srikanth, Bobin Deng, Thomas M. Conte
School of Computer Science
Georgia Institute of Technology

Abstract

The idea of computational error correction has been around for over half a century. The motivation has largely been to mitigate unreliable devices, manufacturing defects or harsh environments, primarily as a mandatory measure to preserve reliability, or more recently, as a means to lower energy by allowing soft errors to occasionally creep. While residue codes have shown great promise for this purpose, there have been several orthogonal non-residue based techniques. In this article, we provide a high level outline of some of these non-residual approaches.

Overview

We first classify various approaches to computational error correction into two broad categories:

1. *Temporal Redundancy.* This approach is based on the hypothesis that the probability of transient errors that occur at the same place to have temporal multiplicity is very low. In other words, a soft error occurs infrequently at the same device, and as such, repeated *measurements* in some manner would serve as an indicator to the correct computation.

2. *Spatial Redundancy.* This approach is based on the hypothesis that the probability of multiple identical computations to all be in error at the same time is very low. In other words, by replicating a computation, any error in a small fraction of the replicas can be *masked* / overpowered by the other correct replicas.

These principles, it turns out, are fundamental to any sort of error correction including computation (ex. arithmetic), storage (ex. memory) and transmission (ex. networking). Some proposals favor spatial redundancy over temporal redundancy, some vice versa, and some employ both, depending upon the target fault model and environment. Given a technique, it is relatively straightforward to determine presence of temporal and/or spatial redundancy, as such, we leave this to the interested reader.

Von Neumann [1] was among the first to propose using redundant components to overcome the effects of defective devices. He introduced the now widely used technique of Triple Modular Redundancy (TMR), which essentially uses three devices instead of one and uses a majority voter to infer a correct output. To note here is that such a mechanism can correct a single error (meaning that at least two of the three devices are not in error), or detect most double errors (where at least

two devices are in error, and their outputs are non-identical). Even today, variants of such a (single error correct, double error detect) SECDED error model are in use.

Also proposed in his work were R-fold modular redundancy (RMR), which reduces to TMR when R=3; Cascade-TMR, which is essentially a multi-level TMR (ex. 3 sets of TMR modules and an overall voter) and NAND multiplexing. The latter replaces each processing unit with multiplexed units containing N lines for each input and output; the multiplex unit itself has two stages: executive stage (performs the function of the processing unit, in parallel) and restorative stage (reduces degradation caused by the executive stage; acts as a non-linear output amplifier). Years later, Nikolic et al. [2] provide a quantitative comparison of these techniques for fault coverage and area overhead necessary. In addition, they also evaluate a reconfiguration technique that uses Configurable Logic Blocks (CLBs) to create fixed function atomic fault tolerant blocks, and clustering them to achieve global fault tolerance.

While these techniques rank high in terms of robustness, they unfortunately, require a very significant overhead in terms of area and energy. It is intuitive that by trading performance, it should be possible to keep component overhead within reasonable limits to achieve reasonable reliability. Arithmetic codes are designed to do just that. We borrow the following classification of arithmetic codes from Wakerly[3].

1. *Separate vs Non-separate.*

   (a) *Separate:* The encoding of a datum X is the concatenation of X and a check symbol computed from X by the function C. The encoding of X is given by $f(X) =< C(X), X >$. Further, upon arithmetic transformation, the data and check symbols are non-interacting. For example, $f(X+Y) =< C(X) * C(Y), X+Y >$.

   (b) *Non-separate:* The encoding of a datum X is simply $C(X)$, with arithmetic transformation being obtained by a single binary operation on the encoding. For example, $f(X+Y) = C(X) * C(Y)$.

2. *Systematic vs Non-systematic.* For systematic codes, the encoding of a datum X has a subfield which equals X.

3. *Homomorphic vs Non-homomorphic.* For homomorphic codes, the data and check symbols are transformed using the same function.

In general, non-separate codes are generally non-systematic (although exception, i.e., systematic non-separate codes exist [29]), and that separate codes are also generally homomorphic (again, exceptions exist [3], however, multiplication is not supported and issues with implementation exist).

A typical example of a separate code is a residue code. In fact, it has been observed that all separate codes are equivalent to residue codes [30]. However, we omit details regarding these as residue codes are not the focus of this article.

The remainder of this article is organized as follows. We first introduce a class of codes known as AN codes, then move onto summarize a large class of various parity predicting / circuit implementation dependent techniques and finally cover other orthogonal approaches to computational error correction.

AN Codes

A typical example of non-separate codes is the AN code [4]. The check function constitutes a simple multiplication by a "check base" A:

$f(X) = AX$; $f(X + Y) = AX + AY$

As such, the code is valid under addition and subtraction, with possible error detection/correction. The crux of such an algorithm relies on the observation that the sum/difference is also a multiple of A. Error detection is relatively straightforward as it just has to verify that the sum is divisible by A. However, error correction involves division by A and a series of subsequent subtractions and bit shifts, the count of which depends upon the number in question; resulting in a multi-variable-cycle latency for error correction, rendering it difficult to design efficient computers around. Years later, Liu [5] proposed a multi-bit error correction for the AN code, albeit at a much higher cost.

Failure to support multiplication notwithstanding, AN codes also run the danger of silent data corruptions due to the inherent possibility of undetected errors (for example, any erroneous number can pass as correct if it happens to be a multiple of A). For these two reasons, proposals to augment AN codes have been made. Forin [6] introduced static signatures (B) to augment the encoding as follows:

$f(X) = AX + B_X$, $0 < B_X < A$ can be arbitrarily chosen for each X.

As such, upon addition of numbers X and Y, the sum modulo A should be $B_X + B_Y$. To detect use of potentially stale (although correct) registers, a timestamping mechanism was further augmented, known as ANBD encoding:

$f(X) = AX + B_X + D$, where D indicates *version*.

Needless to say, for both ANB and ANBD coded systems, software support can be leveraged to hasten error detection by static assignments of signatures and pre-computation of their summations for stack variables. Further, recent work [7] extends this idea to employ signatures at the basic block level to verify dynamic control flow (including non-external function calls) wherever possible, by having instructions communicate counter values to a watchdog entity.

Yet other approaches suggest efficient encoding for multiplication and compiler techniques [8, 9]. However, no error correction scheme is proposed and fail-stop is the default mode of operation.

Micro-architectural / ISA independent Techniques

A different category of 'codes' exists in that they are circuit dependent. Their main idea is to *predict* parity transformation with arithmetic operations and/or careful addition of spatio-temporal redundancy. This constitutes a relatively large body of work, and we will strive to provide an intuitive treatment and coverage of such approaches next.

Sun et al. [10] combine spatial redundancy and temporal redundancy to propose an error correcting parallel adder. They realize a Kogge Stone prefix tree using two Han-Carlson trees, and simply duplicate the generate/propagate circuitry. For elements off the resulting critical path, temporal redundancy is utilized. As far as adders are concerned, their approach is efficient as it is reported to correct 93.76% of soft errors with an area overhead of 12.23% and a delay overhead of 6.41%. However, no such scheme is presented by them for multipliers.

We now review the cost of several competing approaches in chronological order.

Johnson et al. [11] examine exploitation of spatial, temporal and hybrid techniques for the purposes of detection. They note that duplication with comparison incurs over a 100% area overhead. In certain cases, using alternating logic (using compliments) is more efficient *w.r.t.* hardware cost,

however, making a function self-dual may require a 100% increase in area. Finally, for the purposes of temporal redundancy, recomputing with shifted/swapped operands involve negligible area overhead (extra hardware needed only to compute input encoding and output comparison). However, the latter two approaches incur over a 100% increase in latency. While these early implementation seem to be high-cost, especially given that no correction is performed, the ideas are fundamental and are refined over the years.

Hsu et al. [12] propose a temporal redundancy technique that utilizes spatial redundancy in a clever manner. TMR is emulated but hardware overhead is relatively lowered by using 2 x $\frac{n}{3}$ adders/multipliers and using temporal redundancy when an error is detected, to achieve correction. Their approach incurs a hardware overhead of 25% and a delay penalty of 108%.

Nikolaidis [13] propose efficient parity prediction techniques to achieve (detection only) low area overhead of 17% for carry lookahead adders. As noted by the residue based detection work of Pan et al. [14], Nikolaidis et al. propose [15] using differential logic circuits to implement each cell of array-based multipliers, and, also propose [16] output duplicated Booth multipliers, again, for detection alone. The latter was improved upon by Marienfeld et al. [17] to achieve a hardware overhead of 35% for detection in 32 bit multipliers.

Peng et al. [18] develop a mechanism wherein their adder stops upon error detection, and using a deadlock detector, reconfigure the adder. They are able to handle single faults in their 32 bit adders with an 81% area overhead, 140% for 2 faults and 211% for 3 faults. Vasudevan et al. [19] develop error detection in a carry select adder, with an area overhead of 20%.

Rao et al. [20] propose exploiting inherently redundant computation paths in carry generation blocks of carry lookahead adders to identify a faulty block. For the generator and propagator circuitry, time redundancy is used via rotated operands. Ghosh et al. [21] apply temporal redundancy to the Kogge-Stone adder based on the observation that even and odd carries are independent. In their two cycle addition technique, first, one of the correct set of bits (even/odd) are computed and stored at output register. Second, operands are shifted by one bit and the remaining sets of bits (odd/even) are computed and stored.

Rao et al. [22] further propose spatial redundancy for specific parallel prefix adders to achieve fault tolerance. Their design results in area overheads of 85%, 90% and 63% for Brent-Kung, Kogge-Stone and their hybrid implementations respectively.

Valinataj et al. [24] distribute TMR to protect carries alone, the premise being that correct carries are sufficient to generate correct output parities for both addition and multiplication. For a 32 bit carry lookahead adder and a 32 bit wallace tree multiplier, their technique requires an area overhead of 115% and 240% respectively. Krekhov et al. [23] propose parity prediction for the purposes of multi-bit error correction for addition, with roughly a 100% area overhead. Mathew et al. [26] propose parity prediction for multi-bit error detection and correction in Galois Field multipliers with over 100% of area overhead.

Keren et al. [25] observe that not all 2^k outputs of a k-bit output are generally valid outputs for a given combinational circuit. Instead of using redundancy bits, the input to the checker is created using the output bits and the input bits, assuming the function unit is implemented as two independent circuits. Based on the combinational circuits implemented on an FPGA by them, an average overhead of 85% in the number of LUTs was observed.

Dolev et al. [27] seek to transform Hamming codes with arithmetic, by generating codes for the fundamental NAND operation. The idea is to re-generate the code by performing correction on the input and then performing NAND and XOR. The XOR could be replaced with a BCH encoder for multi-bit error correction support.

Other Orthogonal Techniques

Banerjee et al. [28] outline an ASIC design flow where-in the application designer specifies which modules are critical and which can do with approximate outputs. This is especially relevant in DSP applications, for example, in FIR filters, certain coefficients are more important than the others. Post identification via a probabilistic analysis to determine potential fault locations, series transistors are added to mitigate potential shorts and parallel transistors are added to mitigate potential opens.

Blome et al. [33] propose using a small protected cache of live register values for better coverage than protected register files in that protecting the non-state logic (like read/write logic) of a storage structure is easier for smaller structures. As far as computation is concerned, time delayed shadow latches are used to leverage temporal redundancy.

Another class of computational error correction techniques is that of limiting redundancy to that of detection and using rollback to an error-free checkpoint to recover. Needless to mention, such a mechanism comes with its own set of trade-offs and challenges, some of which are: non-zero error detection latency, recovery latency, degree and stride of checkpoint placement etc.. We refer the interested reader to the techniques presented by Habkhi et al. [34] as a starting point for a discussion of these issues and related work in the area. Of further note is the use of checkpoint based recovery in speculative processors.

Conclusion

The need for computational error correction itself has been around for over half a century, with goals ranging from harsh environments, manufacturing defects, intermittently reliable devices and/or near threshold computing. We note several methods to computational error correction in this article, and observe that these are overall relatively less efficient (in that they incur more area overhead and/or latency penalty) when compared with residue codes. To wit, prior work [32, 31] on redundant residue codes achieve computational error correction for addition, subtraction as well as multiplication in an elegant manner with a little over 50% overhead in area and with comparable performance to their non-redundant equivalent. While this is relatively superior to the techniques discussed in this article, it is an open research question as to whether better codes/mechanisms (residue based or otherwise) exist for computational error correction.

References

[1] J. von Neumann, Probabilistic Logics and the Synthesis of Reliable Organisms from Unreliable Components, Auroniaru Studies, C.E. Shannon and J. McCarthy, eds., 43-98, Princeton University Press, Princeton N.J. 1955.

[2] Nikolic, K; Sadek, A; Forshaw, M; Architectures for reliable computing with unreliable nanodevices, Proceedings of the 2001 1st IEEE Conference on Nanotechnology

[3] John Wakerly, Error Detecting Codes, Self-Checking Circuits and Applications, Computer Design and Architecture Series, 1978

[4] Brown, D. T. Error Detecting and Error Correcting Binary Codes for Arithmetic Operations, IRE Trans. Electronic Computers EC-9, 333-337, 1960

[5] Chao-kai Liu, Error-correcting-codes in computer arithmetic, Ph.D. dissertation, University of Illinois at Urbana-Champaign, 1972

[6] Forin, P. Vital coded microprocessor: Principles and application for various transit systems. Proc. IFACGCCT (2014): 79-84.

[7] Schiffel, Ute, et al. ANB-and ANBDmem-encoding: detecting hardware errors in software. International Conference on Computer Safety, Reliability, and Security. Springer Berlin Heidelberg, 2010.

[8] Wappler, Ute, and Christof Fetzer. Hardware failure virtualization via software encoded processing. 2007 5th IEEE International Conference on Industrial Informatics. Vol. 2. IEEE, 2007.

[9] Fetzer, Christof, Ute Schiffel, and Martin Süßkraut. AN-encoding compiler: Building safety-critical systems with commodity hardware. International Conference on Computer Safety, Reliability, and Security. Springer Berlin Heidelberg, 2009.

[10] Sun, Yan, et al. Cost effective soft error mitigation for parallel adders by exploiting inherent redundancy. 2010 IEEE International Conference on Integrated Circuit Design and Technology. IEEE, 2010.

[11] Johnson, Barry W., James H. Aylor, and Haytham H. Hana. "Efficient use of time and hardware redundancy for concurrent error detection in a 32-bit VLSI adder." IEEE journal of solid-state circuits 23.1 (1988): 208-215.

[12] Hsu, Yuang-Ming, and E. E. Swartzlander. "Time redundant error correcting adders and multipliers." Defect and Fault Tolerance in VLSI Systems, 1992. Proceedings., 1992 IEEE International Workshop on. IEEE,

[13] Nicolaidis, Michael. "Carry checking/parity prediction adders and ALUs." IEEE Transactions on Very Large Scale Integration (VLSI) Systems 11.1 (2003): 121-128.

[14] Pan, Abhisek, James W. Tschanz, and Sandip Kundu. "A low cost scheme for reducing silent data corruption in large arithmetic circuits." 2008 IEEE International Symposium on Defect and Fault Tolerance of VLSI Systems. IEEE, 2008.

[15] Nicolaidis, Michael, and Hakim Bederr. "Efficient implementations of self-checking multiply and divide arrays." Proceedings. The European Design and Test Conference EDAC, The European Conference on Design Automation ETC European Test Conference EUROASIC, The European Event in ASIC Design Cat. No. 94TH0634-6. IEEE Comput. Soc. Press, Los Alamitos, CA, USA, 1994.

[16] Nicolaidis, Michael, and R. O. Duarte. "Design of fault-secure parity-prediction booth multipliers." date. Vol. 98. 1998.

[17] Marienfeld, Daniel, et al. "New self-checking output-duplicated booth multiplier with high fault coverage for soft errors." 14th Asian Test Symposium (ATS'05). IEEE, 2005.

[18] Peng, Song, and Rajit Manohar. "Fault tolerant asynchronous adder through dynamic self-reconfiguration." 2005 International Conference on Computer Design. IEEE, 2005.

[19] Vasudevan, Dilip P., and Parag K. Lala. "A technique for modular design of self-checking carry-select adder." 20th IEEE International Symposium on Defect and Fault Tolerance in VLSI Systems (DFT'05). IEEE, 2005.

[20] Rao, Wenjing, Alex Orailoglu, and Ramesh Karri. "Fault identification in reconfigurable carry lookahead adders targeting nanoelectronic fabrics." Eleventh IEEE European Test Symposium (ETS'06). IEEE, 2006.

[21] Ghosh, Swaroop, Patrick Ndai, and Kaushik Roy. "A novel low overhead fault tolerant Kogge-Stone adder using adaptive clocking." Proceedings of the conference on Design, automation and test in Europe. ACM, 2008.

[22] Rao, Wenjing, and Alex Orailoglu. "Towards fault tolerant parallel prefix adders in nanoelectronic systems." 2008 Design, Automation and Test in Europe. IEEE, 2008.

[23] Krekhov, E. V., et al. "A method of monitoring execution of arithmetic operations on computers in computerized monitoring and measuring systems." Measurement Techniques 51.3 (2008): 237-241.

[24] Valinataj, Mojtaba, and Saeed Safari. "Fault tolerant arithmetic operations with multiple error detection and correction." 22nd IEEE International Symposium on Defect and Fault-Tolerance in VLSI Systems (DFT 2007). IEEE, 2007.

[25] Keren, Osnat, et al. "Arbitrary error detection in combinational circuits by using partitioning." 2008 IEEE International Symposium on Defect and Fault Tolerance of VLSI Systems. IEEE, 2008.

[26] Mathew, J., et al. "Multiple bit error detection and correction in GF arithmetic circuits." Electronic System Design (ISED), 2010 International Symposium on. IEEE, 2010.

[27] Dolev, Shlomi, et al. "Preserving Hamming Distance in Arithmetic and Logical Operations." Journal of Electronic Testing 29.6 (2013): 903-907.

[28] Banerjee, Nilanjan, Charles Augustine, and Kaushik Roy. "Fault-tolerance with graceful degradation in quality: A design methodology and its application to digital signal processing systems." 2008 IEEE International Symposium on Defect and Fault Tolerance of VLSI Systems. IEEE, 2008.

[29] Garner, Harvey L. "Error codes for arithmetic operations." IEEE Transactions on Electronic Computers 5 (1966): 763-770.

[30] Peterson, W. W. "On checking an adder." IBM Journal of Research and Development 2.2 (1958): 166-168.

[31] B. Deng, S. Srikanth, E. R. Hein, P. G. Rabbat, T. M. Conte, E. DeBenedictis, J. Cook, Computationally-Redundant Energy-Efficient Processing for Y'all (CREEPY), IEEE International Conference on Rebooting Computing 2016.

[32] R. W. Watson and C. W. Hastings, Self-checked computation using residue arithmetic, Proc. IEEE, vol. 54, pp. 1920-1931, Dec. 1966.

[33] Blome, Jason A., et al. "Cost-efficient soft error protection for embedded microprocessors." Proceedings of the 2006 international conference on Compilers, architecture and synthesis for embedded systems. ACM, 2006.

[34] Tabkhi, Hamed, Seyed Ghassem Miremadi, and Alireza Ejlali. "An asymmetric checkpointing and rollback error recovery scheme for embedded processors." 2008 IEEE International Symposium on Defect and Fault Tolerance of VLSI Systems. IEEE, 2008.